arXiv:1611.08687v2 [cs.DS] 29 Nov 2016

Whom to befriend to influence people

Gennaro Cordasco* Luisa Gargano† Manuel Lafond‡ Lata Narayanan§
Adele A. Rescigno† Ugo Vaccaro† Kangkang Wu §

Abstract

Alice wants to join a new social network, and influence its members to adopt a new product or idea. Each person v in the network has a certain threshold $t(v)$ for *activation*, i.e adoption of the product or idea. If v has at least $t(v)$ activated neighbors, then v will also become activated. If Alice wants to activate the entire social network, whom should she befriend? More generally, we study the problem of finding the minimum number of links that a set of external influencers should form to people in the network, in order to activate the entire social network. This *Minimum Links* Problem has applications in viral marketing and the study of epidemics. Its solution can be quite different from the related and widely studied Target Set Selection problem. We prove that the Minimum Links problem cannot be approximated to within a ratio of $O(2^{\log^{1-\varepsilon} n})$, for any fixed $\varepsilon > 0$, unless $NP \subseteq DTIME(n^{polylog(n)})$, where n is the number of nodes in the network. On the positive side, we give linear time algorithms to solve the problem for trees, cycles, and cliques, for any given set of external influencers, and give precise bounds on the number of links needed. For general graphs, we design a polynomial time algorithm to compute size-efficient link sets that can activate the entire graph.

1 Introduction

The increasing popularity and proliferation of large online social networks, together with the availability of enormous amounts of data about customer bases, has contributed to the rise of *viral marketing* as an effective strategy in promoting new products or ideas. This strategy relies on the insight that once a certain fraction of a social network adopts a product, a larger cascade of further adoptions is predictable due to the *word-of-mouth* network effect [26, 35, 5]. Inspired by social networks and viral marketing, Domingos and Richardson [21, 43] were the first to raise the following important algorithmic problem in the context of social network analysis: If a company can turn a subset of customers in a given network into early adopters, and the goal is to trigger a large cascade of further adoptions, which set of customers should they target?

We use the well-known threshold model to study the influence diffusion process in social networks from an algorithmic perspective. The social network is modelled by a node-weighted graph

*Department of Psychology, Università degli Studi della Campania "Luigi Vanvitelli", Italy
†Department of Computer Science, Università degli Studi di Salerno, Italy
‡Department of Computer Science and Operational Research, Université de Montréal, Montréal, Canada
§Department of Computer Science and Software Engineering, Concordia University, Montréal, Canada

$G = (V, E, t)$ with $V(G)$ representing individuals in the social network, $E(G)$ denoting the social connections, and t an integer-valued *threshold function*. Starting with a *target set*, that is, a subset $S \subseteq V$ of nodes in the graph, that are *activated* by some external incentive, influence propagates deterministically in discrete time steps, and *activates* nodes.

For any unactivated node v, if the number of its activated neighbors at time step $t-1$ is at least $t(v)$, then node v will be activated in step t. A node once activated stays activated. It is easy to see that if S is non-empty, then the process terminates after at most $|V|-1$ steps. We call the set of nodes that are activated when the process terminates as the *activated set*. The problem proposed by Domingo and Richardson [21, 43] can now be formulated as follows: Given a social network $G = (V, E, t)$, and an integer k, find a subset $S \subseteq V$ of size k so that the resulting activated set is as large as possible. In the context of viral marketing, the parameter k corresponds to the budget, and S is a target set that maximizes the size of the activated set. One question of interest is to find the cheapest way to activate the *entire network*, when possible. The optimization problem that results has been called the *Target Set Selection Problem*, and has been widely studied (see for eg. [8, 3, 40]): the goal is to find a minimum-sized set $S \subseteq V$ that activates the entire network (if such a set exists). In a certain sense, the elements of this minimum target set S are the most influential people in the network; if they are activated, the entire network will eventually be activated.

There are, however, two hidden flaws in the formulation of the target set problem. First, the nodes in the target set are assumed to be activated immediately by external incentives, *regardless of their own thresholds of activation*. This is not a realistic assumption; in the context of viral marketing, it is possible, perhaps even likely, that highly influential nodes have high thresholds, and cannot be activated by external incentives alone. Secondly, there is no possibility of giving *partial* external incentives; indeed the target set is activated *only* by external incentives, and the remaining nodes *only* by the internal network effect.

In this paper, we address the flaws mentioned above. We study a related but different problem. Suppose Alice wants to join a new social network, whom should she befriend if her goal is to influence the entire social network? In other words, to whom should Alice create links, so that she can activate the entire network? If Alice creates a link to a node v, the threshold of v is only effectively reduced by one, and so v in turn is activated only if its threshold is one. The problem can be generalized to any set of k external influencers that wish to collectively "take over" a network. We call our problem the *Minimum Links* problem (Min-Links).

The Min-Links problem provides a new way to model a viral marketing strategy, which addresses the flaws described in the target set problem formulation. The links added from the external nodes correspond to the external incentive given to the endpoints of these links. The nodes that are the endpoints of these new links may not be immediately completely activated, but their thresholds are effectively reduced; this corresponds to their receiving partial incentives. One way of seeing this is that every individual to whom we link is given a $10 coupon; for some people this may be enough for them to buy the product, for others, it reduces their resistance to buying it. Individuals with high thresholds cannot be activated only by external incentives. The Min-Links problem also has important applications in epidemiology or the spread of epidemics: in the spread of a new disease, where an infected person or a set of infected people arrives from *outside* a community, the Min-Links problem corresponds to identifying the smallest set of people such that if the infected external people have contact with this set, the entire community could potentially be infected.

Observe that the solution to the Min-Links problem can be quite different from the solution to the Target Set Selection problem for a given network. For example, consider a star network, where

the leaves all have threshold 1, while the central node has degree $|V|-1$ and has threshold $|V|$. The optimal target set is the central node, while the only solution to the Min-Links problem with a single influencer is to create links to all nodes in the network. Thus, a solution to the Min-Links problem can be arbitrarily larger than one to the Target Set Selection problem for the same social network. However, any solution to the Min-Links problem is clearly also a feasible solution to the Target Set Selection problem.

1.1 Our Results

We prove that there exists a (gap-preserving) reduction from the classical Target Set Selection problem to the Min-Links problem. Using the important results by [8], this implies that the Min-Links problem, even in presence of a single external influencer, cannot be approximated to within a ratio of $O(2^{\log^{1-\varepsilon} n})$, for any fixed $\varepsilon > 0$, unless $NP \subseteq DTIME(n^{polylog(n)})$, where n is the number of nodes in the graph. In light of this hardness result, we study the complexity of the problem for networks that can be represented as trees, cycles, and cliques. In each case, we give a necessary and sufficient condition for the feasibility of the Min-Links problem, based on the structural properties and an observation of the threshold function. We then give $O(|V|)$ algorithms to solve the Min-Links problem for all the studied graph topologies. We also give exact bounds on the number of links needed to activate the entire network for all the above specific topologies, as a function of the threshold values. Finally, we present a polynomial time algorithm that, given an arbitrary network G and a number of influencers equal to the maximum node threshold, computes a "small" set of links sufficient to activate the whole network. Our polynomial time algorithm always returns a solution for G of size at most $\sum_{v \in V} \frac{t(v)(t(v)+1)}{2(d_G(v)+1)}$, where $d_G(v)$ is the degree of the vertex v.

1.2 Related work

The problem of identifying the most influential nodes in a social network has received a tremendous amount of attention [27, 31, 37, 29, 9, 28, 4, 23]. The algorithmic question of choosing the target set of size k that activates the most number of nodes in the context of viral marketing was first posed by Domingos and Richardson [21]. Kempe *et al.* [32] started the study of this problem as a discrete optimization problem, and studied it in both the probabilistic independent cascade model and the threshold model of the influence diffusion process. They showed the NP-hardness of the problem in both models, and showed that a natural greedy strategy has a $(1-1/e-\varepsilon)$-approximation guarantee in both models; these results were generalized to a more general cascade model in [33].

In the Target Set Selection problem, the size of the target set is not specified in advance, but the goal is to activate the entire network. Namely, given a graph G and fixed arbitrary thresholds $t(v)$, $\forall v \in V$, find a target set of minimum size that eventually activates all (or a fixed fraction of) nodes of G. Chen [8] proved a strong inapproximability result for the Target Set Selection problem that makes unlikely the existence of an algorithm with approximation factor better than $O(2^{\log^{1-\varepsilon}|V|})$. A polynomial-time algorithm for trees was given in the same paper. Chen's inapproximability result stimulated a series of papers (see for instance [1, 2, 3, 6, 7, 11, 12, 13, 14, 15, 16, 24, 25, 30, 36, 39, 40, 42, 44, 45] and references therein quoted) that isolated many interesting scenarios in which the problem and variants thereof become tractable. Ben-Zwi et al. [3] generalized Chen's result

on trees to show that target set selection can be solved in $n^{O(w)}$ time where w is the treewidth of the input graph. The effect of several parameters, such as diameter and vertex cover number, of the input graph on the complexity of the problem are studied in [40]. The Minimum Target Set has also been studied from the point of view of the spread of disease or epidemics. For example, in [22], the case when all nodes have a threshold k is studied; the authors showed that the problem is NP-complete for fixed $k \geq 3$.

Maximizing the number of nodes activated within a specified number of rounds has also been studied [20, 38]. The problem of dynamos or dynamic monopolies in graphs (eg. [41]) is essentially the target set problem restricted to the case when every node's threshold is half its degree. The recent monograph [10] contains an excellent overview of the area.

The paper closest to our work is [19], in which Demaine *et al.* introduce a model to *partially incentivize* nodes to maximize the spread of influence. Our work differs from theirs in several ways. First, they study the maximization of influence given a fixed budget, while we study in a sense the budget (number of links) needed to activate the entire network. Second, they consider thresholds chosen uniformly at random, while we study arbitrary thresholds. Finally, they allow arbitrary fractional influence to be applied externally on any node, while in our model, every node that receives a link has its threshold reduced by the same amount.

2 Notation and preliminaries.

Given a social network represented by an undirected graph $G = (V, E, t)$, we introduce a set of external nodes U that are assumed to be already activated. We assume that all edges have unit weight; this is generally called the *uniform weight assumption*, and has previously been considered in many papers [8, 25, 14]. A *link set* for (G, U) is a set S of links between nodes in U and nodes in V, i.e $S \subseteq \{(u, v) \mid u \in U; v \in V\}$. For a link set S, we define $E(S) = \{v \in V \mid \exists (u, v) \in S\}$, that is, $E(S)$ is the set of V-endpoints of links in S. For a node v, define $s(v)$ to be the number of links in S for which v is an endpoint. Since the set of external nodes U is already activated, observe that adding the link set S to G is equivalent to reducing the threshold of the node v by $s(v)$. In the viral marketing scenario, the link set S represents giving v a partial incentive of $s(v)$ [18, 19].

Given a link set S for a graph G, we define $I(G, S)$ to be the set of nodes in G that are eventually activated as a result of adding the link set S, that is, by reducing the threshold of each node $v \in E(S)$ by $min\{s(v), t(v)\}$, and then running the influence diffusion process. See Figure 1 for an illustration. Observe that in the target set formulation, this is the same as the set of nodes activated by using U as the target set in the graph G', the graph obtained from G by adding the set U to the node set and the set S to the set of edges.

A link set S such that $I(G, S) = V$, that is, S activates the entire network, is called a *pervading link set*. A pervading link set of minimum size is called an *optimal pervading link set.*

Definition 1 Minimum Links (Min-Links**) problem:** *Given a social network $G = (V, E, t)$, where t is the threshold function on V, and a set of external nodes U, find an optimal pervading link set for (G, U).*

For each node $v \in E(S)$, we say we *give* v a link, or that v *receives* a link. In our algorithms, we express a link set as a *link vector*, $\mathbf{s} = (s(v_1), \ldots, s(v_n))$, where $s(v)$, as defined earlier, is an integer representing the number of links between external nodes in U and the vertex $v \in V$. The

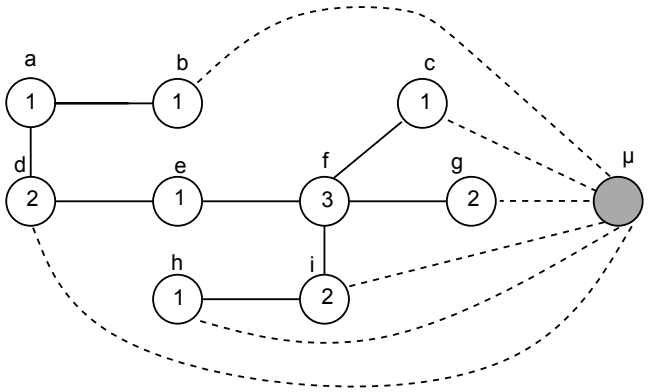

Figure 1: Node μ is the only external influencer and is assumed to be activated. Links in the link set are shown with dashed edges. The given link set activates the entire network and is an optimal pervading link set.

external influencer-endpoints of these links are understood to be distinct, but otherwise can be chosen arbitrarily within U. If activating $X \subseteq V$ activates, directly or indirectly, the set of vertices Y, we write $X \sim Y$ (note that there may be vertices outside Y that X activates). We write $x \sim Y$ instead of $\{x\} \sim Y$. The minimum cardinality of a link set for a Min-Links instance G is denoted $ML(G)$.

Observe that for some graphs, and some sizes of the external influencer set, a pervading link set may not exist. For example, consider a singleton node of threshold greater than 1, and a single influencer. The existence of a feasible solution can be verified in $O(E)$ time by giving k links to every node in V, and simulating the influence diffusion process.

3 NP-hardness

In this section, we consider the complexity of the Min-Links problem. We prove the following result.

Theorem 1 *In networks with n nodes* Min-Links *problem cannot be approximated to within a ratio of* $O(2^{\log^{1-\varepsilon} n})$ *for any fixed* $\varepsilon > 0$*, unless* $NP \subseteq DTIME(n^{polylog(n)})$*, even if the network has bounded degree, and all thresholds are at most 2.*

Proof. We construct a gap-preserving reduction from the Target Set Selection (TSS) problem. We recall that, given an input graph G and a threshold function t, the TSS problem asks for a minimum size subset of vertices of G that can activate all the other vertices. The inapproximability claim of the theorem follows from the inapproximability result of TSS proved in [8], which holds even for graphs with bounded degree, when the thresholds are at most 2.

Starting from an input instance of the TSS problem, that is, a bounded-degree graph $G = (V,E,t)$ with threshold function t, such that $t(v) \leq 2$ for all vertices $v \in V$, we build an instance of the Min-Links problem with a single influencer. Define the graph $G' = (V',E',t')$ as follows:

- $V' = \bigcup_{v \in V} V'_v$ where $V'_v = \{v', v'', v_1, \ldots v_{t(v)}\}$. In particular,

– we replace each $v \in V$ by the gadget Λ_v (cf. Fig. 2) in which the node set is V'_v and v' and v'' are connected by the disjoint paths (v', v_i, v'') for $i = 1, \ldots, t(v)$;

– the threshold of v' in G' is equal to the threshold $t(v)$ of v in G, while each other node in V'_v has threshold equal to 1.

• $E' = \{(v', u') \mid (v,u) \in E\} \cup \bigcup_{v \in V} \{(v', v_i), (v_i, v''), \text{ for } i = 1, \ldots, t(v)\}$.

Summarizing, G' is constructed in such a way that for each gadget Λ_v, the node v' plays the role of v and is connected to all the gadgets representing neighbors of v in G. Hence, G corresponds to the subgraph of G' induced by the set $\{v' \in V'_v \mid v \in V\}$. It is worth mentioning that during an influence diffusion process if any node that belongs to a gadget Λ_v is active, then all the vertices in Λ_v will be activated within the next 3 steps. Moreover, there is only one influencer and the influencer set is $U = \{\mu\}$. Observe that all thresholds in G' are at most 2, and G' remains of bounded degree.

We claim that there is a target set $T \subseteq V$ for G of cardinality $|T| = k$ *if and only* if there is a pervading link set for (G', U) of size k. Assume that $T \subseteq V$ is a target set for G, we consider the set of links S', with $|S'| = k$, defined as

$$S' = \left\{(\mu, v'') \mid v'' \text{ is the extremal node in the gadget } \Lambda_v \text{ and } v \in T\right\}.$$

To see that S' is a pervading link set, we notice that $S' \sim \{u \mid u \in V'_v, v \in T\}$ within three steps. Consequently, recalling that T is a target set and that G is isomorphic to the subgraph of G' induced by $\{v' \in V'_v \mid v \in V\}$, all the vertices $v \in V'$ will be activated, that is $I(G', S') = V'$.

On the other hand, assume that S' is a pervading link set for (G', U) and $|S'| = k$, we can easily build a target set

$$T = \{v \in V \mid \text{there exists } w \in V'_v \text{ such that } (\mu, w) \in S'\}.$$

By construction $|T| \leq |S'|$. We show now that T is a target set for G. To this aim, for each $v \in V$ we consider two cases according to how the node $v' \in \Lambda_v$ associated with v is activated in G':

- If there exists $w \in V'_v$ such that $(\mu, w) \in S'$ then, by construction $v \in T$.
- Suppose otherwise that for each $w \in V'_v$ it holds that $(\mu, w) \notin S'$. In order to activate v' (and afterwards any other node in Λ_v), there must exist a step i when at least $t(v)$ of the neighbors of v' in $V' - V'_v$ are active.

Now we recall that G is the subgraph of G' induced by the set $\{v' \in V'_v \mid v \in V\}$. Hence, for each step $i \geq 0$ and for each v' which is active in G' at step i (with link set S'), we conclude that the corresponding node v must be active in G by step i (with target set T). Consequently any node v will be activated in G. □

In the case of very small degree bound, it has been proved in [34] that the Min-Links problem is NP-hard; in fact, it is almost as hard as Set-Cover to approximate, even if G has degree bounded by 3 and thresholds bounded by 2.

Theorem 2 *[34] The decision version of* Min-Links *is NP-complete, even when restricted to instances with maximum degree* 3 *and maximum threshold* 2*. Moreover, there exists a constant* $\varepsilon > 0$ *such that the optimization version of* Min-Links*, under the same restrictions, is NP-hard to approximate within a* $\varepsilon \ln n$ *factor, where* n *is the number of nodes of the given graph.*

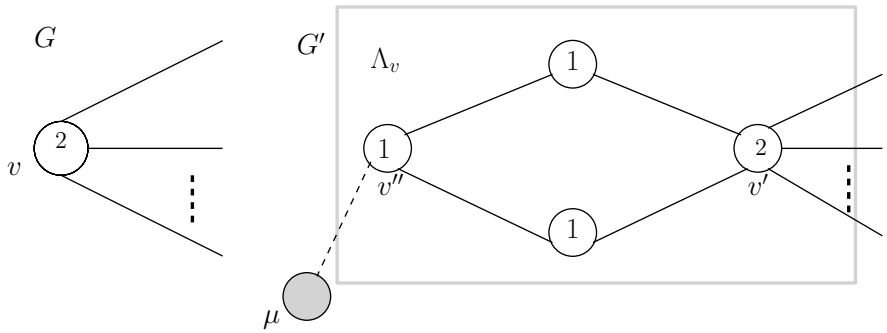

Figure 2: The gadget Λ_v: (left) a generic node $v \in V$ having degree $d_G(v)$ and threshold $t(v) = 2$; (right) the gadget Λ_v, having $t(v) + 2 = 4$ vertices, associated to v.

4 Algorithms for MinLinks

In this section, we give linear time algorithms to solve the MinLinks problem for trees, cycles, and cliques, for any given set of k external influencers (i.e. a single node is able to receive up to k links). Hereafter, the external influencers are $U = \{\mu_1, \ldots, \mu_k\}$ and a solution S for a graph G consists in a set of distinct links (μ_i, v) where $\mu_i \in U$ and $v \in V(G)$. We start with the following simple observation:

Observation 1 *A graph G does not have a pervading link set if it has a node v such that $t(v) > degree(v) + k$, or if every node has threshold strictly greater than k.*

4.1 Trees

In contrast to the NP-completeness of the Min-Links problem shown in the previous section, we now show that there is a linear time algorithm to solve the problem in trees. We start with a necessary and sufficient condition for a tree T to have a valid pervading link set.

Proposition 1 *Let T be a tree and let v be a leaf in T. Let $T' = T - \{v\}$ and T'' be the same as T' except that the threshold of w, the neighbor of v in T, is reduced by 1. Then T has a pervading link set if and only if (a) either $t(v) \leq k$ and T'' has a pervading link set or (b) $t(v) = k+1$ and T' has a pervading link set.*

We now prove a critical lemma that shows that for any node v in the tree, there is an optimal solution that gives $\min(t(v), k)$ links to v.

Lemma 1 *Let T be a tree with n nodes that has a pervading link set, and let v be a node in T. Then there exists an optimal solution for* Min-Links(T) *in which v gets $\min\{t(v), k\}$ links.*

Proof. We prove the lemma by induction on the number of nodes n in the tree. Clearly it is true if $n = 1$. Suppose $n > 1$, and let S be an optimal pervading link set for T. Moreover, choose S such that v receives a maximum number of links among all optimal solutions. If v gets $\min\{t(v), k\}$ links, we are done. If not, then v cannot be activated by external influence alone, and so v must

have a neighbor w that is activated before it, and that contributes to the activation of v. Let T_1 and T_2 be the two trees created by removing the edge between v and w, with T_1 containing w, and let S_1 (respectively S_2) be the links of S with an endpoint in T_1 (respectively T_2). Since T is a tree, and v is activated after w by S, none of the links in S_2 can contribute to the activation of nodes in T_1. It follows that S_1 is a pervading link set for T_1, and in fact is optimal, as a smaller solution for T_1 could be combined with S_2 to yield a better solution for T, contradicting the optimality of S. By the inductive hypothesis, there is an optimal solution S' for T_1 that gives $\min\{t(w),k\}$ links to w. Note that $|S'| = |S_1|$, and $S' \cup S_2$ must also be an optimal solution for T. Let μ_i be an external influencer not giving a link to v in $S' \cup S_2$, and let μ_j be an external influencer giving a link to w in $S' \cup S_2$ (note that μ_i and μ_j must exist). Clearly $S'' = S' \cup S_2 \cup \{(\mu_i, v)\} - \{(\mu_j, w)\}$ also activates the entire tree T (because the w influence on v is replaced by (μ_i, v), and so v still activates, and the (μ_j, w) influence on w is replaced by v's activation). Moreover since $|S''| = |S|$, we conclude that S'' is an optimal solution for T. But S'' gives more links to v than S, contradicting our choice of S. We deduce that there is an optimal pervading link set that gives $\min\{t(v),k\}$ links to v, as needed to complete the proof by induction. □

The above lemma suggests a simple way to break up the Min-Links problem for a tree into subproblems that can be solved independently, which yields a linear-time greedy algorithm.

Theorem 3 *The* Min-Links *problem can be solved for trees in linear time.*

Proof. Given a tree T, let v be an arbitrary leaf in the tree. By Lemma 1, there is an optimal solution, say S, to the Min-Links problem for T that gives $\min\{t(v),k\}$ links to v. Let S_v be the set of links given to v. Suppose $t(v) > k$, then the links given to v are not enough to activate v, and therefore v's neighbor w must contribute to the activation of v. Also, v's activation cannot help in activating any other nodes in T. Thus $S - S_v$ must be an optimal solution to $T' = T - \{v\}$. Suppose instead that $t(v) \leq k$. Then the links given to v activate it immediately. Consider the induced subgraph $T^{(1)}$ of T that contains v, plus every node of T of threshold 1. Let C be the connected component (subtree) of $T^{(1)}$ that contains v (note that C might have only v). Then clearly $v \sim C$. Since S is optimal, S cannot contain any link to a node in C except for v. Construct T' by removing C from T, and subtracting 1 from the threshold of any node x who is a neighbor of a node in C. Observe that any such node x can be a neighbor of exactly one node in C, since T is a tree. Then $S - S_v$ must be an optimal solution to T'; if instead there is a smaller-sized solution to T', we can add the links from S_v to v to that solution to obtain a smaller solution for T than S, contradicting the optimality of S.

The above argument justifies the correctness of the following simple greedy algorithm. Initialize $S = \emptyset$. Take a leaf v in the tree. If $t(v) > k+1$ then there is no solution by Observation 1. If $t(v) = k+1$, then give k links to v in S, remove v from the tree, and recursively solve the remaining tree. If $t(v) \leq k$, then give $t(v)$ links to v (from arbitrary influencers), remove the subtree of T that is connected to v consisting only of nodes of degree 1, reduce the thresholds of all neighbors of the nodes in this subtree by 1, and recursively solve the resulting trees. It is easy to see that the algorithm can be implemented in linear time. □

For the network in Figure 1, assuming that leaves in the tree are always processed in alphabetical order, the greedy algorithm given in Theorem 3 first picks node b and adds a link to it. We then remove nodes b and a, and reduce the threshold of d by 1. Next we pick c, give it a link, remove it from the tree, and decrement $t(f)$ to 2. The next leaf that is picked and given a link is d; since d's

threshold now is 1, we remove d and e from the tree, and reduce f's threshold to 1. Proceeding in this way, we arrive at the link set shown.

We now give an exact bound on $ML(T)$, the number of links required to activate the entire tree T:

Theorem 4 *Let T be a tree that has a pervading link set. Then $ML(T) = 1 + \sum_{v \in T}(t(v) - 1)$*

Proof. We give a proof by induction on the number of nodes n in the tree. Clearly if the tree consists of a single node x, there is a solution if and only if $t(x) \leq k$, and the number of links needed is $t(x)$ which is equal to $1 + \sum_{v \in V}(t(v) - 1)$ as needed. Now consider a tree T with $n > 1$ nodes and let x be a leaf in the tree. Then by Lemma 1, there is an optimal solution S in which x gets a set S_x of $\min\{t(x), k\}$ links. By Observation 1, there is a solution only if $t(x) \leq k+1$. Let $T' = T - \{x\}$ (all nodes keep the same thresholds as in T) and let T'' be the tree derived from T by removing x and reducing the threshold of w, the neighbor of x in T by 1.

First we consider the case when $t(x) = k+1$. Then giving k links to x from S_x is not sufficient to activate it. By the usual cut-and-paste argument, $S - S_x$ must be an optimal solution for tree T'.

$$\begin{aligned} ML(T) &= k + ML(T') \\ &= t(x) - 1 + (1 + \sum_{v \in V(T')} (t(v) - 1)) \text{ by the inductive hypothesis} \\ &= 1 + \sum_{v \in V(T)} (t(v) - 1). \end{aligned}$$

Next we consider the case when $t(x) \leq k$, and $t(w) > 1$. Then x is immediately activated by the $t(x)$ links it receives in S, and the activation of x effectively reduces the threshold of w. Therefore, $S - S_x$ must be an optimal solution for the tree T'' in which the threshold of w is $t(w) - 1$. It follows that

$$\begin{aligned} ML(T) &= t(x) + ML(T'') \\ &= t(x) + (1 + \sum_{v \in V(T'')} (t(v) - 1)) \text{ by the inductive hypothesis} \\ &= t(x) + 1 + (t(w) - 2) + \sum_{v \in V(T'') - \{w\}} (t(v) - 1) \\ &= 1 + \sum_{v \in V(T)} (t(v) - 1). \end{aligned}$$

Finally suppose $t(x) \leq k$ and $t(w) = 1$. Then it is impossible that S contains a link to w, as this would contradict the optimality of S. Therefore, we can move one link from node v to node w, to get a new optimal pervading link set S' for T. Furthermore, $S' - S_x$ must also be an optimal pervading link set for T'. It follows that

$$\begin{aligned} ML(T) &= ML(T') \\ &= t(x) - 1 + (1 + \sum_{v \in V(T')} (t(v) - 1)) \text{ by the inductive hypothesis} \\ &= 1 + \sum_{v \in V(T)} (t(v) - 1). \end{aligned}$$

□

We remark that in contrast to the intuition for the optimal target set problem, where we would choose nodes of high degree or threshold to be in the target set, in the Min-Links problem, our algorithm gives links to leaves initially, though eventually nodes that were internal nodes in the tree may also receive links. That is, the best nodes to befriend might be the nodes with a single connection to other nodes in the tree!

4.2 Cycles

In this section, we give a solution for the Min-Links problem on cycles. Let $C_n = (V, E, t)$ be a cycle with n nodes, $V = \{0, 1, ..., n-1\}$, $E = \{((i, i+1) \ mod \ n) \mid 1 \leq i \leq n\}$, and $t : t(v) \to \mathcal{Z}^+$. We define $P_{i,j}$ $(i \neq j)$ to be the sub-path of C_n consisting of all nodes in $\{i, \ldots, j\}$ in the clockwise direction. We may use the $[i, j]$ notation to denote the vertices of $P_{i,j}$. By *consecutive vertices of threshold* $k+2$, we mean two vertices i, j such that the only two vertices in $P_{i,j}$ with threshold $k+2$ are i and j.

Proposition 2 *A cycle has a pervading link set if and only if the following conditions hold:*
(1) there is at least one node of threshold at most k,
(2) every node is of threshold at most $k+2$,
(3) between any two consecutive nodes of threshold $k+2$*, there is at least one node of threshold at most k.*

Proof. The necessity of the first two conditions follows from Observation 1. As for the third condition, suppose there are two consecutive nodes i and j of threshold $k+2$, such that all nodes between them have threshold $k+1$. Then both nodes i and j needs both their neighbors to be activated before them (even if they receive k links), but meanwhile, since there is no node of threshold k or less in $[i+1, j-1]$, no node in the sub-path $P_{i+1,j-1}$ can be activated. Therefore none of the nodes in the sub-path $P_{i,j}$ can be activated. Conversely, if all three conditions listed in the statement are met, it is easy to see that by giving k links to every node in the cycle, all the nodes in the cycle can be activated. □

We note that a similar condition can be stated for paths, with the additional restriction that there must be a node of threshold at most k before (after) the first (last resp.) node of threshold $k+2$, if any.

We give a linear time algorithm for finding a minimum-sized link set for problem Min-Links(C_n). Essentially we reduce the problem to finding an optimal solution for an appropriate path.

Theorem 5 *The* Min-Links *problem for a cycle* C_n *can be solved in linear time.*

Proof. If all nodes are of threshold 1, or if there is a single node with threshold 2, and the remaining nodes all have threshold 1, then by giving a link to *any* of the nodes with threshold 1, we can activate the entire cycle, and this is clearly optimal.

Therefore, in what follows, we assume that one of the following cases holds:

(a) the minimum threshold is greater than 1,

(b) the minimum threshold is 1, and there are at least two nodes with threshold ≥ 2, or

(c) the minimum threshold is 1 and there is exactly one node with threshold > 2.

Fix an arbitrary node i of minimum threshold in C_n. We define $c(i)$ and $cc(i)$ to be the first node with threshold > 1 in i's clockwise direction and counter clockwise direction respectively. Observe that in cases (a) and (b), $c(i) \neq cc(i)$ (see Figure 3), and in case (c) above, $c(i) = cc(i)$. We also define $P_{c(i),cc(i)}$ to be the path from $c(i)$ to $cc(i)$, except that in cases (a) and (b), we decrement $t(c(i))$ and $t(cc(i))$ by 1; and in case (c), the path contains a single node $c(i)$, and we decrement the threshold of $c(i)$ by 2. We now prove that an optimal solution to MinLinks(C_n) can be constructed by giving $t(i)$ links to i and combining them with an optimal solution to $P_{c(i),cc(i)}$.

We first claim that there exists an optimal solution that gives $t(i)$ links to i. To see this, let S be an optimal solution that gives $q < t(i)$ links to node i. Observe that $q \in \{t(i)-1, t(i)-2\}$, as otherwise it is impossible for i to be activated. First suppose $q = t(i)-1$. This means one of i's neighbours must activate i. We follow the chain of activation to i, which must start at some node j which is activated entirely by external influence. That is, S must give $t(j)$ links to j. Without loss of generality, we assume $j \sim j+1 \sim \ldots \sim i-1 \sim i$. Let μ be such that $(\mu, j) \in S$ and $(\mu, i) \notin S$. Such a μ must exist since j received $t(j)$ links and $t(j) \geq t(i)$ as i was a node of minimum threshold in C_n while i received $< t(i)$ links by assumption. We now construct a new solution $S' = S - \{(\mu, j)\} \cup \{(\mu, i)\}$ by moving a link from j to i. Since i receives $t(i)$ links in S', the node i is immediately activated. Furthermore, $i \sim i-1 \sim \ldots \sim j+1$. Finally, since j receives $t(j)-1$ links in S', and $j+1$ is activated, node j is activated in the next step. Thus S' is a solution of the same size as S, that gives $t(i)$ links to i as needed.

Next suppose $q = t(i)-2$. Then both neighbours of i must be activated by S before i, and both serve to activate i. We then follow the chain of activation in both directions from i and find nodes p and q that were activated entirely by external influence (it is possible that $p = q$). We then move a link from each of p and q to i. Now i is activated entirely by external influence and eventually activates p and q. This completes the proof of the claim that there exists an optimal solution that gives $t(i)$ links to i.

Consider therefore an optimal solution S that gives $t(i)$ links to the node i; let us call this set of links S_i. It is not hard to see that that $S - S_i$ must be an optimal solution to Min-Links($P_{c(i),cc(i)}$), since activating i activates $[cc(i)+1, c(i)-1]$ and lowers the threshold of $cc(i)$ and $c(i)$ by 1 each in cases (a) and (b) above and lowers the threshold of $c(i)$ by 2 in case (c) above.

Finally, since the Min-Links problem for a path can be solved in linear time according to Theorem 3, we can construct an optimal solution for a cycle in linear time as well. □

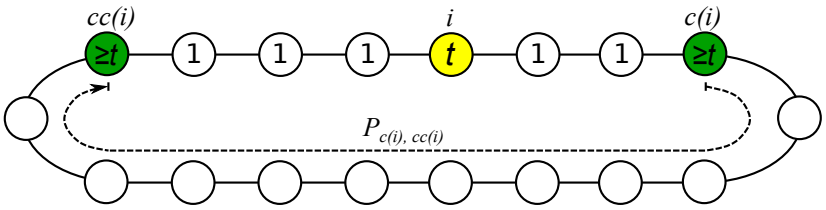

Figure 3: A cycle in which case (b) holds. Node i has minimum threshold t, while $cc(i)$ and $c(i)$ are the nodes closest to i that have threshold higher than 1.

We give an exact bound on the number of links required to fully activate a cycle.

Theorem 6 *Given a cycle $C_n = (V,E,t)$ which has a pervading link set, then*

$$ML(C_n) = max\{1, \sum_{j=1}^{n}(t(j)-1)\}.$$

Proof. If all nodes have threshold 1, then $ML(C_n) = 1 = max\{1, \sum_{j=1}^{n}(t(j)-1)\}$. If one node has threshold 2, and all the remaining nodes have threshold 1, then $ML(C_n) = 1 = \sum_{j=1}^{n}(t(j)-1)$. Finally, for all remaining cases, it follows from the optimality of our algorithm that $ML(C_n) = t(i) + ML(P_{cc(i),c(i)})$ where i is a node of minimum threshold in C_n, and the value of $t(c(i)) + t(cc(i))$ is 2 less in $P_{cc(i),c(i)}$ than in C_n. By Theorem 4, we have $ML(P_{cc(i),c(i)}) = -1 + \sum_{j\in[cc(i),c(i)]}(t(j)-1)$. Therefore $ML(C_n) = t(i) - 1 + \sum_{j\in[cc(i),c(i)]}(t(j)-1) = \sum_{j=1}^{n}(t(j)-1)$. □

4.3 Cliques

In this section, we give an algorithm to solve the Min-Links problem on cliques. Let $K_n = (V,E,t)$ be a clique with n nodes, $V = \{1,2,...,n\}$ and $E = \{(i,j) : 1 \le i < j \le n\}$ and $t : t(v) \to \mathcal{Z}^+$. We first show a necessary and sufficient condition for the Min-Links problem to have a feasible solution:

Proposition 3 *Let K_n be a clique with $t(i) \le t(i+1)$, for all $1 \le i < n$. Then K_n has a pervading link set if and only if $t(i) \le i+k-1$ for all $1 \le i \le n$.*

Proof. If $t(i) \le k+i-1$ for all $1 \le i \le n$, it is easy to see that there exists a solution S by giving k links to every node i; we claim that node i is activated in or before round i. Since $t(1) \le k$, node 1 is activated in round 1. Inductively, node 1 to $i-1$ are already activated in round $i-1$, the effective threshold of node i has been reduced to at most k. Node i receives k links, therefore, node i must be activated in the i^{th} round, if it is not already activated. Conversely, suppose there exist nodes j such that $t(j) > k+j-1$ and there exists a solution S to the Min-Links problem; let p be the smallest such node with $t(p) > k+p-1$. In order to activate any node $q \ge p$, at least p nodes have to be activated before q, since $t(q) \ge t(p) > k+p-1$. However, there are only $p-1$ nodes that can be activated before any such node $q \ge p$. Thus no node q with $q \ge p$ can be activated, a contradiction. □

We now give a greedy algorithm to solve the Min-Links problem on a clique.

Theorem 7 *The* Min-Links *problem for a clique K_n can be solved in time $\Theta(n+k)$.*

Proof. First sort the nodes in order of threshold. By Observation 1, there is no solution if any node has a threshold greater than $n+k$, therefore, we can use counting sort and complete the sorting in $\Theta(n+k)$ time. Clearly, the condition given in Proposition 3 can easily be checked in linear time. We now give the following greedy linear time algorithm for a clique which has a feasible solution: give $t(1)$ links to node 1 in order to activate it, and let $j > 1$ be the smallest value such that node j is not activated after activating node 1. Remove all nodes in $\{1,\ldots,j-1\}$, decrement by $j-1$ the thresholds of all nodes $\ge j$, and solve the resulting graph recursively. It is easy to see that this algorithm can be implemented in linear time, in an iterative fashion as follows: we examine the

nodes in order. When we process node i, if $t(i) < i$, we simply increment i and continue; if $t(i) \geq i$, we give $t(i) - i + 1$ links to node i (note that we assume $t(i) \leq k + i - 1$, and so $t(i) - i + 1 \leq k$). We now show that the link set produced by this greedy algorithm is optimal.

First we show that there must be an optimal solution that gives $t(1)$ links to the node 1. Consider an optimal solution S in which node 1 gets $q < t(1)$ links. We follow the chain of activations by S to node 1. Let A_1 be the set of nodes that is activated by external influence alone according to solution S, and suppose $A_1 \sim A_2 \sim \ldots \sim A_j \sim 1$. Then $|A_1 \cup A_2 \cup \cdots \cup A_j| \geq t(1) - q$. We construct a new solution S' by moving $t(1) - q$ links from some node $i \in A_1$ to node 1. Observe that i received $t(i) \geq t(1)$ links in S, and so there are enough links to move. Furthermore, the nodes in $A'_1 = A_1 - \{i\} \cup \{1\}$ are activated entirely by external influence in S'. Also since $|A'_1| = |A_1|$, and all nodes are connected, $A'_1 \sim A_2 \ldots \sim A_j$. Finally, together with the $t(i) - t(1) + q$ links that S' gives to i, we have $A_1 \cup A_2 \cup A_j \sim i$. The rest of the activation proceeds in the same way in S and S'. Since S' is a solution of the same size as S that gives $t(1)$ links to node 1, the claim is proved.

Let $S_1 \subseteq S$ be the links given to node 1 in S. Next we claim that $S - S_1$ is an optimal solution to the clique C' which is the induced sub-graph on the nodes $\{j, j+1, \ldots, n\}$ where $j > 1$ is the smallest index with $t(j) \geq j$, and with thresholds of all nodes reduced by $j - 1$. Suppose there is a smaller solution S' to C'. We claim that $S' \cup S_1$ activates all nodes in the clique K_n. Since for any node $1 < k < j$, we have $t(k) < k$, it can be seen inductively that S_1 suffices to activate node k. Thus, all nodes in $\{1, 2, \ldots j-1\}$ are activated. Furthermore, the thresholds of all nodes in $\{j, j+1, \ldots, n\}$ are effectively reduced by $j-1$. Thus using the links in S' suffices to activate them. Finally, since $|S'| < |S - S_1|$, $S' \cup S_1$ is a smaller solution than S to the clique K_n, contradicting the optimality of S for K_n. We conclude that the greedy algorithm described above produces a minimum sized solution to the Min-Links problem. □

The following tight bound on the minimum number of links to activate an entire clique is immediate:

Theorem 8 *Given a clique K_n which has a feasible solution, let $P = \{i : t(i) \geq i\}$. Then*

$$ML(K_n) = \sum_{i \in P} t(i) - i + 1.$$

5 An algorithm for general graphs

In this section we design an algorithm, that works for arbitrary graphs $G = (V, E)$, to efficiently allocate links to nodes in V from a set of external nodes U in such a way that, assuming that the nodes in U are already activated, they trigger an influence diffusion process that activates the whole network. We assume $|U| \geq t_{max}$, where t_{max} is the maximum threshold of nodes in V so that there is always an feasible solution for the Min-Links problem. Our procedure is formally presented in Algorithm 1.

The algorithm works by computing a link vector $\mathbf{s} = (s(v_1), \ldots, s(v_n))$, where $s(v)$ is an integer representing the number of links between external nodes in U and the vertex $v \in V$. As observed in Section 2, the $s(v)$ links to v can be seen as a partial incentive to node v. From any link vector $\mathbf{s}$, one can get link sets S between nodes in U and nodes in V. In the following we say that a link vector $\mathbf{s}$ is a pervading link vector when any corresponding set S is a *pervading link set*.

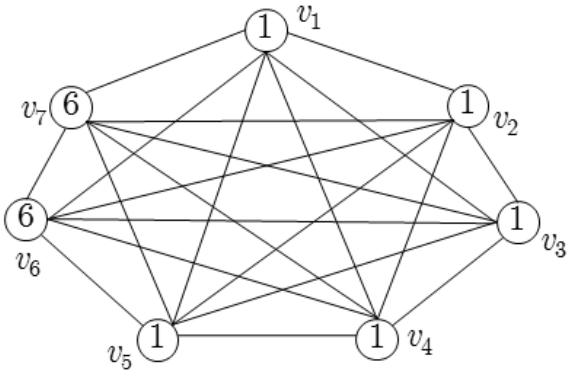

Figure 4: An example to illustrate Algorithm TPI(G). The number inside each circle is the node threshold.

The algorithm proceeds by iteratively deleting nodes from the graph G, and at each iteration the node to be deleted is chosen to maximize a certain parameter (Case 2). If, during the deletion process, a node v in the surviving graph remains with less neighbors than its current threshold (Case 1), then a set of links (or equivalently a partial incentive) is added to v so that v's new threshold is equal to the number of neighbours of v in the surviving graph.

In the sequel, we denote by $I(G,\mathbf{s},j)$ the set of nodes that are active at step j of the influence diffusion process on the network G augmented with a set of links determined by the link vector $\mathbf{s}$. Namely, let $\Gamma_G(v)$ denote the neighborhood of v in G, we have $I(G,\mathbf{s},0)=\{v \mid s(v)\geq t(v)\}$ and $I(G,\mathbf{s},j)=I(G,\mathbf{s},j-1)\cup\{u \mid |\Gamma_G(v)\cap I(G,\mathbf{s},j-1)\geq t(u)-s(u)\}$, for all $j\geq 1$. The pseudocode for our algorithm is given in Algorithm TPI(G).

Example 1 *Consider a complete graph on 7 nodes with thresholds $t(v_1)=\ldots=t(v_5)=1$, $t(v_6)=t(v_7)=6$ (cf. Fig. 4). A possible execution of the algorithm is summarized below. At each iteration of the while loop, the algorithm considers the nodes in the order shown in the table below, where we also indicate for each node whether Cases 1 or 2 applies and the updated value of the number of links for the selected node:*

Iteration	1	2	3	4	5	6	7	8
Node	v_7	v_6	v_6	v_1	v_2	v_3	v_4	v_5
Case	2	1	2	2	2	2	2	1
Links	0	1	1	0	0	0	0	1

The algorithm $TPI(G)$ outputs the link vector $\mathbf{s}$ having non zero elements $s(v_5)=s(v_6)=1$, for which we have

$$\begin{aligned}
I(G,\mathbf{s},0) &= \{v_5\} \qquad (\text{since } s(v_5)=1=t(v_5))\\
I(G,\mathbf{s},1) &= I(G,\mathbf{s},0)\cup\{v_1,v_2,v_3,v_4\}=\{v_1,v_2,v_3,v_4,v_5\}\\
I(G,\mathbf{s},2) &= I(G,\mathbf{s},1)\cup\{v_6\}=\{v_1,v_2,v_3,v_4,v_5,v_6\} \qquad (\text{since } s(v_6)=1)\\
I(G,\mathbf{s},3) &= I(G,\mathbf{s},2)\cup\{v_7\}=V.
\end{aligned}$$

Algorithm 1: Algorithm TPI(G)

Input: A graph $G = (V, E, t)$ where t is the threshold function on V.

Output: $\mathbf{s} = (s(v_1), \ldots, s(v_n))$ a link vector for G, where $V = \{v_1, \ldots, v_n\}$.

$W = V$;
foreach $v \in V$ **do**
    $s(v) = 0$; // # of links to v.
    $\delta(v) = d(v)$;
    $k(v) = t(v)$;
    $N(v) = \Gamma_G(v)$;
while $W \neq \emptyset$ **do** // Select one node and either update its links or remove it from the graph.
    if $\exists\, v \in W$ *s.t.* $k(v) > \delta(v)$ **then** // Case 1: Increase $s(v)$ and update $k(v)$.
        $s(v) = s(v) + k(v) - \delta(v)$;
        $k(v) = \delta(v)$;
        if *k(v)=0* **then** // here $\delta(v) = 0$.
            $W = W - \{v\}$;
    else // Case 2: Choose a node v to eliminate from the graph.
        $v = \texttt{argmax}_{u \in W} \left\{ \frac{k(u)(k(u)+1)}{\delta(u)(\delta(u)+1)} \right\}$;
        foreach $u \in N(v)$ **do**
            $\delta(u) = \delta(u) - 1$;
            $N(u) = N(u) - \{v\}$;
        $W = W - \{v\}$;
    return s

We first prove the correctness and complexity of the algorithm. Subsequently, we compute an upper bound on the number of links in the pervarding link set, that is, an upper bound on $\sum_{v\in V} s(v)$. To this end, we define the following notation.

Let ℓ be the number of iterations of the while loop in TPI(G). For each iteration j of the while loop, with $1 \leq j \leq \ell$, we denote

- by W_j the set W at the beginning of the j-th iteration (cf. line 7 of $TPI(G)$), in particular $W_1 = V(G)$ and $W_{\ell+1} = \emptyset$;
- by $\mathcal{G}(j)$ the subgraph of G induced by the vertices in W_j;
- by v_j the node selected during the j-th iteration[1];
- by $\delta_j(v)$ the degree of node v in $\mathcal{G}(j)$;
- by $k_j(v)$ the value of the threshold of node v in $\mathcal{G}(j)$, that is, as it is updated at the beginning of the j-th iteration, in particular $k_1(v) = t(v)$ for each $v \in V$;
- by $s_j(v)$ the number of links to v that are computed by the algorithm from the j-th iteration until and including the ℓ-th iteration, in particular observe that $s_{\ell+1}(v) = 0$ and $s_1(v) = s(v)$ for each $v \in V$;
- by σ_j the number of links assigned *during* the j-th iteration, that is,

$$\sigma_j = s_j(v_j) - s_{j+1}(v_j) = \begin{cases} 0 & \text{if } k_j(v_j) \leq \delta_j(v_j), \\ k_j(v_j) - \delta_j(v_j) & \text{otherwise.} \end{cases}$$

According to the above notation, we have that if node v is selected during the iterations $j_1 < j_2 < \ldots < j_{a-1} < j_a$ of the while loop in TPI(G), where the last value j_a is the iteration when v has been eliminated from the graph, then

$$s_j(v) = \begin{cases} \sigma_{j_1} + \sigma_{j_2} + \ldots + \sigma_{j_a} & \text{if } j \leq j_1, \\ \sigma_{j_b} + \sigma_{j_{b+1}} + \ldots + \sigma_{j_a} & \text{if } j_{b-1} < j \leq j_b \leq j_a, \\ 0 & \text{if } j > j_a. \end{cases}$$

In particular observe that for $j = j_a$, we have $s_j(v_j) = \sigma_j$.

The following result is immediate.

Proposition 1 *Consider the node v_j that is selected during iteration j, for $1 \leq j \leq \ell$, of the while loop in the algorithm TPI(G):*

(1.1) *If Case 1 of the algorithm TPI(G) holds and $\delta_j(v_j) = 0$, then $k_j(v_j) > \delta_j(v_j) = 0$ and the isolated node v_j is eliminated from $\mathcal{G}(j)$. Moreover,*

$$W_{j+1} = W_j - \{v_j\}, \quad s_{j+1}(v_j) = s_j(v_j) - \sigma_j, \quad \sigma_j = k_j(v_j) - \delta_j(v_j) = k_j(v_j) > 0,$$

and, for each $v \in W_{j+1}$

$$s_{j+1}(v) = s_j(v), \quad \delta_{j+1}(v) = \delta_j(v), \quad k_{j+1}(v) = k_j(v).$$

[1] A node can be selected several times before being eliminated; indeed in Case 1 we can have $W_{j+1} = W_j$.

(1.2) *If Case 1 of TPI(G) holds with $\delta_j(v_j) > 0$, then $k_j(v_j) > \delta_j(v_j) > 0$ and no node is deleted from $\mathcal{G}(j)$, that is, $W_{j+1} = W_j$. Moreover,*

$$\sigma_j = k_j(v_j) - \delta_j(v_j) > 0$$

and, for each $v \in W_{j+1}$

$$\delta_{j+1}(v) = \delta_j(v), \quad s_{j+1}(v) = \begin{cases} s_j(v_j) - \sigma_j & \text{if } v = v_j \\ s_j(v) & \text{if } v \neq v_j \end{cases}, \quad k_{j+1}(v) = \begin{cases} \delta_j(v) & \text{if } v = v_j \\ k_j(v) & \text{if } v \neq v_j. \end{cases}$$

(2) *If Case 2 of TPI(G) holds then $k_j(v_j) \leq \delta_j(v_j)$ and v_j is pruned from $\mathcal{G}(j)$. Hence,*

$$W_{j+1} = W_j - \{v_j\}, \qquad \sigma_j = 0,$$

and, for each $v \in W_{j+1}$ it holds

$$s_{j+1}(v) = s_j(v), \quad k_{j+1}(v) = k_j(v) \quad \delta_{j+1}(v) = \begin{cases} \delta_j(v) - 1 & \text{if } v \in \Gamma_{\mathcal{G}(j)}(v_j) \\ \delta_j(v) & \text{otherwise.} \end{cases}$$

Lemma 2 *For each iteration $j = 1, 2, \ldots, \ell$, of the while loop in the algorithm TPI(G),*

(1) if $k_j(v_j) > \delta_j(v_j)$ then $\sigma_j = k_j(v_j) - \delta_j(v_j) = 1$;

(2) if $\delta_j(v_j) = 0$ then $s_j(v_j) = k_j(v_j)$.

Proof. First, we prove (1). At the beginning of the algorithm, $t(u) = k(u) \leq d(u) = \delta(u)$ for all $u \in V$. Afterwards, the value of $\delta(u)$ is decreased by at most one unit for each iteration (cf. line 16 of *TPI*(G)). Moreover, the first time the condition of Case 1 holds for some node u, one has $\delta_j(u) = k_j(u) - 1$. Hence, if the selected node is $v_j = u$ then (1) holds; otherwise, some $v_j \neq u$, satisfying the condition of Case 1 is selected and $\delta_{j+1}(u) = \delta_j(u)$ and $k_{j+1}(u) = k_j(u)$ hold. Hence, when at some subsequent iteration $j' > j$ the algorithm selects $v_{j'} = u$, we have $\delta_{j'}(u) = k_{j'}(u) - 1$. To show (2), it is sufficient to notice that at the iteration j when node v_j is eliminated from the graph, $s_j(v) = \sigma_j$. □

We are now ready to prove the correctness and complexity of the algorithm TPI(G):

Theorem 9 *For any graph $G = (V, E, t)$ and a set U of external influencers such that $|U| \geq max_{v \in V} t(v)$, the algorithm TPI(G) outputs a pervading link vector for G in time $O(|E| \log |V|)$.*

Proof. We show that for each iteration j, with $1 \leq j \leq \ell$, the assignment of $s_j(v)$ links for each $v \in W_j$ activates all the nodes of the graph $\mathcal{G}(j)$ when the distribution of thresholds to its nodes is $k_j(\cdot)$. The proof is by induction on j.

If $j = \ell$ then the unique node v_ℓ in $\mathcal{G}(\ell)$ has degree $\delta_\ell(v_\ell) = 0$ and $s_\ell(v_\ell) = k_\ell(v_\ell) = 1$ (see Lemma 2).

Consider now $j < \ell$ and suppose the algorithm is correct on $\mathcal{G}(j+1)$ that is, the assignment of $s_{j+1}(v)$ links to each $v \in W_{j+1}$, activates all the nodes of the graph $\mathcal{G}(j+1)$ when the distribution of thresholds to its nodes is $k_{j+1}(\cdot)$.

Recall that v_j denotes the node the algorithm selects from W_j (thus obtaining W_{j+1}, the node set of $\mathcal{G}(j+1)$). In order to prove the theorem we analyze three cases according to the current degree and threshold of the selected node v_j.

- $k_j(v_j) > \delta_j(v_j) = 0$: By Lemma 2, we have $k_j(v_j) = s_j(v_j)$. Then the correctness of the algorithm on $\mathcal{G}(j)$ follows from part (1.1) of Proposition 1 and the inductive hypothesis on $\mathcal{G}(j+1)$.

- $k_j(v_j) > \delta_j(v_j) \geq 1$: From (1.2) of Proposition 1, we observe that $k_j(v) - s_j(v) = k_{j+1}(v) - s_{j+1}(v)$, for each node $v \in W_j$. Indeed, for each $v \neq v_j$ we have $k_{j+1}(v) = k_j(v)$ and $s_{j+1}(v) = s_j(v)$. Moreover,

$$k_{j+1}(v_j) - s_{j+1}(v_j) = \delta_j(v_j) - s_j(v_j) + \sigma_j = k_j(v_j) - s_j(v_j).$$

  Hence the nodes that can be activated in $\mathcal{G}(j+1)$ can be activated in $\mathcal{G}(j)$ with thresholds $k_j(\cdot)$ and $s_j(\cdot)$ links. So, the correctness of the algorithm on $\mathcal{G}(j)$ follows from the inductive hypothesis on $\mathcal{G}(j+1)$.

- $k_j(v_j) \leq \delta_j(v_j)$: From part (2) of Proposition 1, and by the inductive hypothesis on $\mathcal{G}(j+1)$, we have that all the neighbors of v_j in $\mathcal{G}(j)$ that are nodes in W_{j+1} gets activated; since $k_j(v_j) \leq \delta_j(v_j)$, we conclude that v_j also gets activated in $\mathcal{G}(j)$.

This completes the proof by induction. To prove the time complexity, we first notice that the algorithm ends within $|E|$ iterations of the *while* loop. Indeed, each time a vertex v is selected in Case 1 of the algorithm, its current threshold $k(v)$ is decreased to the current node degree $\delta(v)$. Moreover, if $k(v)$ reaches 0 then v is deleted. This implies that each vertex v can be selected up to $d(v)$ times before being deleted (because its threshold gets down to 0 or because Case 2 is applied to v). Hence, the algorithm executes at most $|E|$ iterations of the while loop.

When Case 2 applies, one has to select the node u in the current node set that maximizes the quantity $\frac{k(u)(k(u)+1)}{\delta(u)(\delta(u)+1)}$. Suppose that the nodes are initially sorted in a max-heap according to the priorites $\frac{k(u)(k(u)+1)}{\delta(u)(\delta(u)+1)}$. Each time the node priority $\frac{k(u)(k(u)+1)}{\delta(u)(\delta(u)+1)}$ changes, we can update the heap in $O(\log|V|)$ time. For each iteration, if v is the processed node then the update involves only the neighbors of v if Case 2 applies, and only the node v itself if Case 1 applies. Overall, Case 1 requires at most $|E| \cdot O(\log|V|)$ time and Case 2 requires at most $\sum_{v\in V} d(v) \cdot O(\log|V|)$ times.

Hence, the algorithm can be implemented in such a way to run in $O(|E|\log|V|)$ time. □

We now compute an upper bound on the number of links that the link vector returned by the algorithm TPI assigns to the vertices in V.

Theorem 10 *For any graph G the algorithm TPI(G) returns a link vector* **s** *for G such that*

$$\sum_{v\in V} s(v) \leq \sum_{v\in V} \frac{t(v)(t(v)+1)}{2(d_G(v)+1)}.$$

Proof. Define $\mathrm{B}(j) = \sum_{v\in W_j} \frac{k_j(v)(k_j(v)+1)}{2(\delta_j(v)+1)}$, for each $j = 1,\ldots,\ell$. By definition of ℓ, we have $\mathcal{G}(\ell+1)$ is the empty graph; we then define $\mathrm{B}(\ell+1) = 0$. We prove now by induction on j that

$$\sigma_j \leq \mathrm{B}(j) - \mathrm{B}(j+1). \tag{1}$$

By using (1) we will have the bound on $\sum_{v\in V} s(v)$. Indeed,

$$\sum_{v\in V} s(v) = \sum_{j=1}^{\ell} \sigma_j \leq \sum_{j=1}^{\ell} (\mathrm{B}(j) - \mathrm{B}(j+1)) = \mathrm{B}(1) - \mathrm{B}(\ell+1) = \mathrm{B}(1) = \sum_{v\in V} \frac{t(v)(t(v)+1)}{2(d(v)+1)}.$$

In order to prove (1), we analyze three cases depending on the relation between $k_j(v_j)$ and $\delta_j(v_j)$.
• Assume first $k_j(v_j) > \delta_j(v_j) = 0$. We get

$$\begin{aligned}
\mathrm{B}(j) - \mathrm{B}(j{+}1) &= \sum_{v\in W_j} \frac{k_j(v)(k_j(v)+1)}{2(\delta_j(v)+1)} - \sum_{v\in W_{j+1}} \frac{k_{j+1}(v)(k_{j+1}(v)+1)}{2(\delta_{j+1}(v)+1)} \\
&= \frac{k_j(v_j)(k_j(v_j)+1)}{2(\delta_j(v_j)+1)} + \sum_{v\in W_j-\{v_j\}} \frac{k_j(v)(k_j(v)+1)}{2(\delta_j(v)+1)} \\
&\qquad - \sum_{v\in W_{j+1}} \frac{k_{j+1}(v)(k_{j+1}(v)+1)}{2(\delta_{j+1}(v)+1)} \\
&= \frac{k_j(v_j)(k_j(v_j)+1)}{2(\delta_j(v_j)+1)} \qquad \text{(by 1.1 in Proposition 1)} \\
&= 1 = \sigma_j. \qquad \text{(by Lemma 2)}
\end{aligned}$$

• Let now $k_j(v_j) > \delta_j(v_j) \geq 1$. We have

$$\begin{aligned}
\mathrm{B}(j) - \mathrm{B}(j{+}1) &= \sum_{v\in W_j} \frac{k_j(v)(k_j(v)+1)}{2(\delta_j(v)+1)} - \sum_{v\in W_{j+1}} \frac{k_{j+1}(v)(k_{j+1}(v)+1)}{2(\delta_{j+1}(v)+1)} \\
&= \frac{k_j(v_j)(k_j(v_j)+1)}{2(\delta_j(v_j)+1)} - \frac{k_{j+1}(v_j)(k_{j+1}(v_j)+1)}{2(\delta_{j+1}(v_j)+1)} \\
&\quad + \sum_{v\in W_j-\{v_j\}} \frac{k_j(v)(k_j(v)+1)}{2(\delta_j(v)+1)} - \sum_{v\in W_{j+1}-\{v_j\}} \frac{(k_{j+1}(v)(k_{j+1}(v)+1)}{2(\delta_{j+1}(v)+1)} \\
&= \frac{(\delta_j(v_j)+1)(\delta_j(v_j)+2)}{2(\delta_j(v_j)+1)} - \frac{\delta_j(v_j)(\delta_j(v_j)+1)}{2(\delta_j(v_j)+1)} \quad \text{(by 1.2 in Proposition 1)} \\
&= \frac{2(\delta_j(v_j)+1)}{2(\delta_j(v_j)+1)} = 1 = \sigma_j. \qquad \text{(by Lemma 2)}
\end{aligned}$$

• Finally, let $k_j(v_j) \leq \delta_j(v_j)$. In this case, by the algorithm we know that

$$\frac{k_j(v)(k_j(v)+1)}{\delta_j(v)(\delta_j(v)+1)} \leq \frac{k_j(v_j)(k_j(v_j)+1)}{\delta_j(v_j)(\delta_j(v_j)+1)}, \tag{2}$$

for each $v \in W_j$. Hence, we get

$$\begin{aligned}
\mathrm{B}(j)-\mathrm{B}(j+1) &= \sum_{v\in W_j}\frac{k_j(v)(k_j(v)+1)}{2(\delta_j(v)+1)} - \sum_{v\in W_{j+1}}\frac{k_{j+1}(v)(k_{j+1}(v)+1)}{2(\delta_{j+1}(v)+1)} \\
&= \frac{k_j(v_j)(k_j(v_j)+1)}{2(\delta_j(v_j)+1)} + \sum_{v\in\Gamma_{\mathcal{G}(j)}(v_j)}\frac{k_j(v)(k_j(v)+1)}{2(\delta_j(v)+1)} \\
&\quad - \sum_{v\in\Gamma_{\mathcal{G}(j)}(v_j)}\frac{k_{j+1}(v)(k_{j+1}(v)+1)}{2(\delta_{j+1}(v)+1)} \qquad \text{(by 2 in Proposition 1)} \\
&= \frac{k_j(v_j)(k_j(v_j)+1)}{2(\delta_j(v_j)+1)} + \sum_{v\in\Gamma_{\mathcal{G}(j)}(v_j)}\frac{k_j(v)(k_j(v)+1)}{2}\left(\frac{1}{(\delta_j(v)+1)}-\frac{1}{\delta_j(v)}\right) \\
&= \frac{k_j(v_j)(k_j(v_j)+1)}{2(\delta_j(v_j)+1)} - \sum_{v\in\Gamma_{\mathcal{G}(j)}(v_j)}\frac{k_j(v)(k_j(v)+1)}{2\delta_j(v)(\delta_j(v)+1)} \\
&\geq \frac{k_j(v_j)(k_j(v_j)+1)}{2(\delta_j(v_j)+1)} - \frac{k_j(v_j)(k_j(v_j)+1)\delta_j(v_j)}{2\delta_j(v_j)(\delta_j(v_j)+1)} \qquad \text{(by (2))} \\
&= 0 = \sigma_j
\end{aligned}$$

□

Experimental data showing the effectiveness of algorithm TPI(G) on real-life networks are given in [18], as well as proofs of its optimality in case of cliques and trees. However, the time complexity does not becomes linear as for the algorithms given in Section 4.

6 Discussion

In this paper, we introduced and studied the Min-Links problem: given a social network G where every node v has a threshold $t(v)$ to be activated, which minimum-sized set of nodes should an already activated set of external influencers S befriend, so as to influence the entire network? We showed that the problem cannot be approximated to within a ratio of $O(2^{\log^{1-\varepsilon} n})$, for any fixed $\varepsilon > 0$, unless $NP \subseteq DTIME(n^{polylog(n)})$. In contrast, we gave exact linear time algorithms that solve the problem in trees, cycles, and cliques, for any given set of k external influencers. We also gave an exact bound (as a function of the thresholds) on the number of links needed for such graphs. Moreover, we gave a polynomial time algorithm that solves the problem in general graphs and derived an upper bound on the number of links used by the algorithm. It would be interesting to generalize these algorithms to find the minimum number of links required to influence a specified fraction of the nodes. Other directions include studying the case with non-uniform weights on the edges. Clearly, the problem remains NP-complete in general, but the complexity for special classes of graphs remains open. Another interesting question is that of maximizing the number of activated nodes, given a fixed budget of ℓ links.

References

[1] E. ACKERMAN, O. BEN-ZWI, G. WOLFOVITZ. Combinatorial model and bounds for target set selection. *Theoretical Computer Science*, Vol. 411, 4017–4022, 2010.

[2] C. BAZGAN, M. CHOPIN, A. NICHTERLEIN and F. SIKORA. Parametrized approximability of maximizing the spread of influence in networks. *Journal of Discrete Algorithms* Vol. 27, 54–65, 2014.

[3] O. BEN-ZWI, D. HERMELIN, D. LOKSHTANOV and I. NEWMAN. Treewidth governs the complexity of target set selection. *Discrete Optimization*, Vol. 8, 87–96, 2011.

[4] C. BORGS, M. BRAUTBAR, J. CHAYES, and B. LUCIER. Maximizing social influence in nearly optimal time. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA '14, 946–957, 2014.

[5] J. JOHNSON BROWN and P. H REINGEN. Social ties and word-of-mouth referral behavior. *Journal of Consumer research*, 350–362, 1987.

[6] C.C. CENTENO M.C. DOURADO, L. DRAQUE PENSO, D. RAUTENBACH and J.L. SZWARCFITER. Irreversible conversion of graphs. *Theoretical Computer Science,* 412 (29), 3693–3700, 2011.

[7] C.-L. CHANG. Triggering cascades on undirected connected graphs. *Information Processing Letters*, 111(19), 973–978, 2011.

[8] N. CHEN. On the approximability of influence in social networks. *SIAM J. Discrete Math.*, 23, 1400–1415, 2009.

[9] W. CHEN, Y WANG and S. YANG. Efficient Influence Maximization in Social Networks. *Proc. 15th ACM SIGKDD Intl. Conf. on Know. Dis. and Data Min.*, 199–208, 2009.

[10] W. CHEN, L, V.S. LAKSHMANAN, and C. CASTILLO. *Information and Influence Propagation in Social Networks*. Morgan & Claypool, 2013.

[11] C.-Y. CHIANG, L.-H. HUANG and H.-G. YEH. Target Set Selection problem for honeycomb networks. *SIAM J. Discrete Math.*, 27(1), 310–328, 2013.

[12] C.-Y. CHIANG, L.-H. HUANG, B.-J. LI, J. WU and H.-G. YEH. Some results on the target set selection problem. *Journal of Combinatorial Optimization*, Vol. 25, 702–715, 2013.

[13] M. CHOPIN, A. NICHTERLEIN, R. NIEDERMEIER and M. WELLER. Constant Thresholds Can Make Target Set Selection Tractable. *Theory of Computing Systems* vol. 55, 61–83, 2014.

[14] F. CICALESE, G. CORDASCO, L. GARGANO, M. MILANIČ, J. G. PETERS and U. VACCARO. Spread of Influence in Weighted Networks under Time and Budget Constraints. *Theoretical Computer Science*, Vol. 586, 40–58, 2015.

[15] F. CICALESE, G. CORDASCO, L. GARGANO, M. MILANIČ and U. VACCARO. Latency-Bounded Target Set Selection in Social Networks. *Theoretical Computer Science*, Vol. 535, 1–15, 2014.

[16] A. COJA-OGHLAN, U. FEIGE, M. KRIVELEVICH and D. REICHMAN. Contagious Sets in Expanders. In: *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, 1953–1987, 2015.

[17] G. CORDASCO, L. GARGANO, M. MECCHIA, A.A. RESCIGNO and U. VACCARO. A fast and effective heuristic for discovering small target sets in social networks. *Proc. of 9th Annual International Conference on Combinatorial Optimization and Applications (COCOA 2015),* Z. Lu *et al.* (Eds) Lecture Notes in Computer Science, volume 9486, 193–208, 2015.

[18] G. CORDASCO, L. GARGANO, A.A RESCIGNO, and U. VACCARO, Optimizing spread of influence in social networks via partial incentives, Proc. of *22nd International Colloquium on Structural Information and Communication Complexity*, Lecture Notes in Computer Science, volume 9439, 119–134, 2015.

[19] E. D. DEMAINE, M. TAGHI HAJIAGHAYI, H. MAHINI, D. L. MALEC, S. RAGHAVAN, A. SAWANT, and M. ZADIMOGHADAM. How to influence people with partial incentives. In *Proceedings of the 23rd International Conference on World Wide Web*, WWW '14, 937–948, 2014.

[20] T. DINH, H. ZHANG, D. NGUYEN, and M. THAI. Cost-effective viral marketing for time-critical campaigns in large-scale social networks. *IEEE/ACM Transactions on Networking*, PP(99):1–1, 2014.

[21] P. DOMINGOS and M. RICHARDSON. Mining the network value of customers. *Proc. of 7th ACM SIGKDD Int. Conf. on Know. Disc. and Data Min.*, 57–66, 2001.

[22] P. DREYER JR. and F. ROBERTS. Irreversible -threshold processes: Graph-theoretical threshold models of the spread of disease and of opinion. *Discrete Applied Mathematics*, 157(7), 1615–1627, 2009.

[23] M. A. FAZLI, M. GHODSI, J. HABIBI, P. JALALY KHALILABADI, V. MIRROKNI, and S. SADEGHABAD. On the non-progressive spread of influence through social networks. In *Latin American Theoretical Informatics Symposium, LATIN 12*, volume 7256 of *Lecture Notes in Computer Science*, 315–326, 2012.

[24] D. FREUND, M. POLOCZEK and D. REICHMAN. Contagious Sets in Dense Graphs in: *Proceedings of 26th Int'l Workshop on Combinatorial Algorithms (IWOCA2015)*, Volume 9538 of *Lecture Notes in Computer Science*, 185–196, 2016.

[25] L. GARGANO, P. HELL, J. G. PETERS and U. VACCARO, Influence Diffusion in Social Networks under Time Window Constraints. *Theoretical Computer Science*, vol. 584, 53–66, 2015.

[26] J. GOLDENBERG, B. LIBAI, and E. MULLER. Talk of the network: A complex systems look at the underlying process of word-of-mouth. *Marketing Letters*, 12, 211–223, 2001.

[27] A. GOYAL, F. BONCHI, L. LAKSHMANAN, and S. VENKATASUBRAMANIAN. On minimizing budget and time in influence propagation over social networks. *Social Network Analysis and Mining*, 3, 179–192, 2013.

[28] A. GOYAL, F. BONCHI, and L. LAKSHMANAN. A data-based approach to social influence maximization. *Proceedings of the VLDB Endowment*, 5, 73–84, 2011.

[29] A. GOYAL, W. LU, and L. LAKSHMANAN. Celf++: Optimizing the greedy algorithm for influence maximization in social networks. In *Proceedings of the 20th International Conference Companion on World Wide Web*, WWW '11, 47–48, 2011.

[30] A. GUGGIOLA and G. SEMERJIAN. Minimal Contagious Sets in Random Regular Graphs. *Journal of Statistical Physics*, 158(2), 300–358, 2015.

[31] J. HE, S. JI, R. BEYAH, and Z. CAI. Minimum-sized influential node set selection for social networks under the independent cascade model. In *Proceedings of the 15th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, MobiHoc '14, 93–102, 2014.

[32] D. KEMPE, J. KLEINBERG, and É. TARDOS. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '03, 137–146, 2003.

[33] D. KEMPE, J. KLEINBERG, and É. TARDOS. Influential nodes in a diffusion model for social networks. In *Proceedings of the 32rd International Conference on Automata, Languages and Programming*, ICALP'05, 1127–1138, 2005.

[34] M. LAFOND, L. NARAYANAN and K. WU. Whom to befriend to influence people. In Proc. of *23rd International Colloquium on Structural Information and Communication Complexity*, Lecture Notes in Computer Science, 2016, to appear.

[35] J. LESKOVEC, L. A ADAMIC, and B. A HUBERMAN. The dynamics of viral marketing. In *Proceedings of the 7th ACM Conference on Electronic Commerce*, 228–237. ACM, 2006.

[36] X. LIU, Z. YANG and W. WANG. Exact solutions for Latency-Bounded Target Set Selection Problem on some special families of graphs. *Discrete Applied Mathematics*, vol. 203: 111–116, 2016.

[37] W. LU, F. BONCHI, A. GOYAL, and L. LAKSHMANAN. The bang for the buck: Fair competitive viral marketing from the host perspective. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '13, 928–936, 2013.

[38] S. LV and L. PAN. Influence maximization in independent cascade model with limited propagation distance. In *Web Technologies and Applications*, volume 8710 of *Lecture Notes in Computer Science*, 23–34, 2014.

[39] F. MORONE and H. A. MAKSE. Influence maximization in complex networks through optimal percolation. *Nature* 524, 65-68, 2015.

[40] A. NICHTERLEIN, R. NIEDERMEIER, J. UHLMANN, and M. WELLER. On tractable cases of target set selection. *Social Network Analysis and Mining*, 3(2), 233–256, 2012.

[41] D. PELEG. Local majority voting, small coalitions and controlling monopolies in graphs: A review. *Theoretical Computer Science*, 282, 231–257, 2002.

[42] T. V. T. REDDY and C. P. RANGAN. Variants of spreading messages. *J. Graph Algorithms Appl.*, 15(5), 683-699, 2011.

[43] M. RICHARDSON and P. DOMINGOS. Mining knowledge-sharing sites for viral marketing. In *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '02, 61–70, 2002.

[44] C. WANG, L. DENG, G. ZHOUB and M. JIANG. A global optimization algorithm for target set selection problems. *Information Sciences* 267(20), 101-118, 2014.

[45] M. ZAKER. On dynamic monopolies of graphs with general thresholds. *Discrete Mathematics*, 312(6), 1136-1143, 2012.